

SPRING Working Group
Internet-Draft
Intended Status: Standards Track
Expires: May 7, 2020

R. Gandhi, Ed.
C. Filsfils
Cisco Systems, Inc.
D. Voyer
Bell Canada
M. Chen
Huawei
B. Janssens
Colt
November 4, 2019

**Performance Measurement Using TWAMP Light
for Segment Routing Networks
draft-gandhi-spring-twamp-srpm-04**

Abstract

Segment Routing (SR) leverages the source routing paradigm. SR is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) data planes. This document specifies procedure for sending and processing probe query and response messages for Performance Measurement (PM) in Segment Routing networks. The procedure uses the mechanisms defined in [RFC 5357](#) (Two-Way Active Measurement Protocol (TWAMP) Light) and Simple Two-Way Active Measurement Protocol (STAMP) for Delay Measurement, and also uses the mechanisms defined in this document for Loss Measurement. The procedure specified is applicable to SR-MPLS and SRv6 data planes and are used for both links and end-to-end SR Policies.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Conventions Used in This Document	4
2.1.	Requirements Language	4
2.2.	Abbreviations	4
2.3.	Reference Topology	5
3.	Overview	5
3.1.	Example Provisioning Model	6
3.2.	STAMP Applicability	7
4.	Probe Messages	7
4.1.	Probe Query Message	7
4.1.1.	Delay Measurement Probe Query Message	7
4.1.1.1.	Delay Measurement Authentication Mode	8
4.1.1.2.	Loss Measurement Probe Query Message	8
4.1.1.2.1.	Loss Measurement Authentication Mode	12
4.1.1.3.	Probe Query for SR Links	13
4.1.1.4.	Probe Query for End-to-end Measurement for SR Policy	13
4.1.1.4.1.	Probe Query Message for SR-MPLS Policy	13
4.1.1.4.2.	Probe Query Message for SRv6 Policy	13
4.2.	Probe Response Message	14
4.2.1.	One-way Measurement Mode	18
4.2.2.	Two-way Measurement Mode	18
4.2.2.1.	Return Path TLV	18
4.2.2.2.	Probe Response Message for SR-MPLS Policy	20
4.2.2.3.	Probe Response Message for SRv6 Policy	20
4.2.3.	Loopback Measurement Mode	21
5.	Performance Measurement for P2MP SR Policies	21
6.	ECMP Support for SR Policies	22
7.	Additional Message Processing Rules	22
7.1.	TTL Value	22
7.2.	Router Alert Option	22
7.3.	UDP Checksum	23

8.	Security Considerations	23
9.	IANA Considerations	23
10.	References	24
10.1.	Normative References	24
10.2.	Informative References	25
	Acknowledgments	27
	Authors' Addresses	27

[1.](#) Introduction

Segment Routing (SR) leverages the source routing paradigm and greatly simplifies network operations for Software Defined Networks (SDNs). SR is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) data planes. SR takes advantage of the Equal-Cost Multipaths (ECMPs) between source and transit nodes, between transit nodes and between transit and destination nodes. SR Policies as defined in [[I-D.spring-segment-routing-policy](#)] are used to steer traffic through a specific, user-defined paths using a stack of Segments. Built-in SR Performance Measurement (PM) is one of the essential requirements to provide Service Level Agreements (SLAs).

The One-Way Active Measurement Protocol (OWAMP) defined in [[RFC4656](#)] and Two-Way Active Measurement Protocol (TWAMP) defined in [[RFC5357](#)] provide capabilities for the measurement of various performance metrics in IP networks using probe messages. These protocols rely on control-channel signaling to establish a test-channel over an UDP path. These protocols lack support for direct-mode Loss Measurement (LM) to detect actual data traffic loss which is required in SR networks. The Simple Two-way Active Measurement Protocol (STAMP) [[I-D.ippm-stamp](#)] alleviates the control-channel signaling by using configuration data model to provision a test-channel. The TWAMP Light [Appendix I in [RFC5357](#)] [[BBF.TR-390](#)] provides simplified mechanisms for active performance measurement in Customer IP networks by provisioning UDP paths and eliminates the control-channel signaling.

This document specifies procedures for sending and processing probe query and response messages for Performance Measurement in SR networks. The procedure uses the mechanisms defined in [[RFC5357](#)] (TWAMP Light) and STAMP for Delay Measurement (DM), and also uses the mechanisms defined in this document for Loss Measurement. The procedure specified is applicable to SR-MPLS and SRv6 data planes and are used for both links and end-to-end SR Policies. This document also defines mechanisms for handling ECMPs of SR Policies for performance delay measurements.

2. Conventions Used in This Document

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2.2. Abbreviations

BSID: Binding Segment ID.

DM: Delay Measurement.

ECMP: Equal Cost Multi-Path.

HMAC: Hashed Message Authentication Code.

LM: Loss Measurement.

MPLS: Multiprotocol Label Switching.

NTP: Network Time Protocol.

OWAMP: One-Way Active Measurement Protocol.

PM: Performance Measurement.

PSID: Path Segment Identifier.

PTP: Precision Time Protocol.

SID: Segment ID.

SL: Segment List.

SR: Segment Routing.

SRH: Segment Routing Header.

SR-MPLS: Segment Routing with MPLS data plane.

SRv6: Segment Routing with IPv6 data plane.


STAMP: Simple Two-way Active Measurement Protocol.

TC: Traffic Class.

TWAMP: Two-Way Active Measurement Protocol.

2.3. Reference Topology

In the reference topology shown below, the sender node R1 initiates a probe query for performance measurement and the responder node R5 sends a probe response for the query message received. The probe response is sent to the sender node R1. The nodes R1 and R5 may be directly connected via a link enabled with Segment Routing or there exists a Point-to-Point (P2P) SR Policy [\[I-D.spring-segment-routing-policy\]](#) on node R1 with destination to node R5. In case of Point-to-Multipoint (P2MP), SR Policy originating from source node R1 may terminate on multiple destination leaf nodes [\[I-D.spring-sr-p2mp-policy\]](#).



Reference Topology

3. Overview

For one-way, two-way and round-trip delay measurements in Segment Routing networks, the TWAMP Light procedures defined in [Appendix I of \[RFC5357\]](#) are used. For one-way and two-way direct-mode and inferred-mode loss measurements in Segment Routing networks, the procedures defined in this document are used. One-way loss measurement provides receive packet loss whereas two-way loss measurement provides both transmit and receive packet loss. Separate UDP destination port numbers are user-configured for delay and loss measurements from the range specified in [\[I-D.ippm-stamp\]](#). The sender uses the UDP port number following the guidelines specified in [Section 6 in \[RFC6335\]](#). For both links and end-to-end SR Policies, no PM session for delay or loss measurement is created on the responder node R5 [\[RFC5357\]](#).

For Performance Measurement, probe query and response messages are sent as following:

- o For Delay Measurement, the probe messages are sent on the congruent path of the data traffic by the sender node, and are

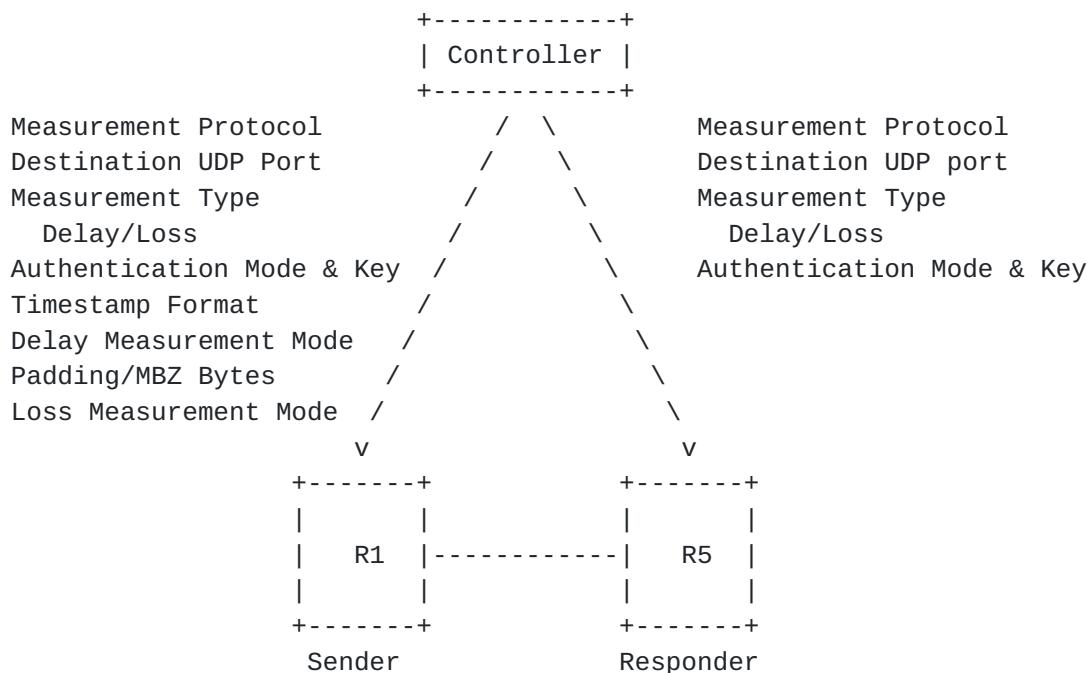
used to measure the delay experienced by the actual data traffic flowing on the links and SR Policies.

- o For Loss Measurement, the probe messages are sent on the congruent path of the data traffic by the sender node, and are used to collect the receive traffic counters for the incoming link or incoming SID where the probe query messages are received at the responder node (incoming link or incoming SID needed since the responder node does not have PM session state present).

The In-Situ Operations, Administration, and Maintenance (IOAM) mechanisms for SR-MPLS defined in [[I-D.spring-ioam-sr-mpls](#)] and for SRv6 defined in [[I-D.spring-ioam-srv6](#)] are used to carry PM information such as timestamp in-band as part of the data packets, and are outside the scope of this document.

3.1. Example Provisioning Model

An example of a provisioning model and typical measurement parameters for performance delay and loss measurements using TWAMP Light is shown in the following Figure:



The responder node R5 uses the parameters for the timestamp format, delay measurement mode, padding size and loss measurement mode from the received probe query message.

The mechanisms to provision the sender and responder nodes are outside the scope of this document.

3.2. STAMP Applicability

The Simple Two-way Active Measurement Protocol (STAMP) [[I-D.ippm-stamp](#)] and the STAMP TLVs [[I-D.ippm-stamp-option-tlv](#)] are both equally applicable to the procedures specified in this document.

This is because the delay measurement message formats defined for STAMP and STAMP TLVs are backwards compatible with the delay measurement message formats defined in [[RFC5357](#)]. Hence, the same user-configured destination UDP port for delay measurement can be used for STAMP and TWAMP Light messages. The STAMP with a TLV for "direct measurement" can be used for combined delay + loss measurement using a separate user-configured UDP destination port.

The loss measurement probe and query messages defined in this document are also equally applicable to STAMP and STAMP TLVs, and use the message formats defined in [[I-D.ippm-stamp](#)].

4. Probe Messages

4.1. Probe Query Message

In this document, the probe messages defined in [[RFC5357](#)] are used for Delay and Loss measurements for SR links and end-to-end SR Policies. The user-configured destination UDP ports (separate UDP ports for different delay and loss message formats) are used for identifying the PM probe packets as described in [Appendix I of \[RFC5357\]](#).

4.1.1. Delay Measurement Probe Query Message

The message content for Delay Measurement probe query message using UDP header [[RFC768](#)] is shown in Figure 1. The DM probe query message is sent with user-configured Destination UDP port number for DM. The Destination UDP port cannot be used as Source port, since the message does not have any indication to distinguish between query and response. The DM probe query message contains the payload for delay measurement defined in [Section 4.1.2 of \[RFC5357\]](#). For symmetrical size query and response messages [[RFC6038](#)], the DM probe query message contains the payload format defined in [Section 4.2.1 of \[RFC5357\]](#).

```

+-----+
| IP Header                                     |
. Source IP Address = Sender IPv4 or IPv6 Address .

```



```

. Destination IP Address = Responder IPv4 or IPv6 Address .
. Protocol = UDP .
. .
+-----+
| UDP Header |
. Source Port = As chosen by Sender .
. Destination Port = User-configured Port for Delay Measurement.
. .
+-----+
| Payload = Message as specified in Section 4.2.1 of RFC 5357 | |
. Payload = Message as specified in Section 4.1.2 of RFC 5357 | .
. Payload = Message specified in Section 4.2 of [I-D.ippm-stamp].
. .
+-----+

```

Figure 1: DM Probe Query Message

Timestamp field is eight bytes and use the format defined in [Section 4.2.1 of \[RFC5357\]](#) . It is recommended to use the IEEE 1588v2 Precision Time Protocol (PTP) truncated 64-bit timestamp format [[IEEE1588](#)] as specified in [[RFC8186](#)].

[4.1.1.1](#). Delay Measurement Authentication Mode

When using the authenticated mode for delay measurement, the matching authentication type (e.g. HMAC-SHA-256) and key are user-configured on both the sender and responder nodes. A separate user-configured destination UDP port is used for the delay measurement in authentication mode due to the different probe message format.

[4.1.1.2](#). Loss Measurement Probe Query Message

In this document, new probe query message formats are defined for loss measurement as shown in Figure 3A and Figure 3B. The message formats are hardware efficient due to the small size payload and well-known locations of counters. They are similar to the delay measurement message formats and do not require any backwards compatibility and support for the existing DM message formats from [[RFC5357](#)].

The message content for Loss Measurement probe query message using UDP header [[RFC768](#)] is shown in Figure 2. The LM probe query message is sent with user-configured Destination UDP port number for LM. Separate Destination UDP ports are used for direct-mode and inferred-mode loss measurements. The Destination UDP port cannot be used as Source port, since the message does not have any indication to distinguish between query and response. The LM probe query message contains the payload for loss measurement as defined in

Figure 3A and Figure 3B. For symmetrical size query and response messages [RFC6038], the LM probe query message contains the payload format defined in Figure 7A and Figure 7B for loss measurement.

```

+-----+
| IP Header                                     |
. Source IP Address = Sender IPv4 or IPv6 Address .
. Destination IP Address = Responder IPv4 or IPv6 Address .
. Protocol = UDP .
. .
+-----+
| UDP Header                                   |
. Source Port = As chosen by Sender .
. Destination Port = User-configured Port for Loss Measurement .
. .
+-----+
| Payload = Message as specified in Figure 3A or 3B |
. Payload = Message as specified in Figure 3C or 3D .
. .
+-----+

```

Figure 2: LM Probe Query Message

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
|                               Sequence Number                               |
+-----+-----+-----+-----+
|                               Transmit Counter                               |
|                               |                                               |
+-----+-----+-----+-----+
|X|B|I|  Reserved                |          Block Number                |
+-----+-----+-----+-----+
|                               Packet Padding                               |
.                               .                                               .
|                               +-----+-----+-----+-----+         |
|                               |          Checksum Complement                |
+-----+-----+-----+-----+

```

Figure 3A: LM Probe Query Message Format

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+

```

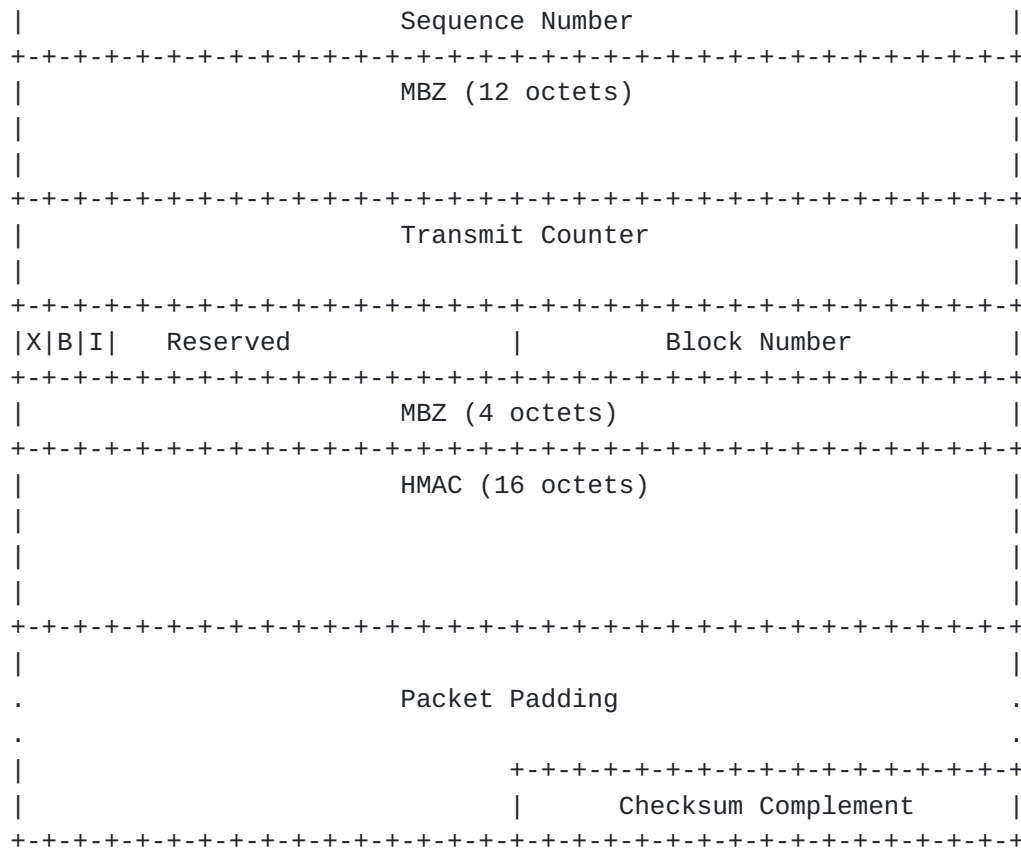



Figure 3B: LM Probe Query Message Format - Authenticated Mode

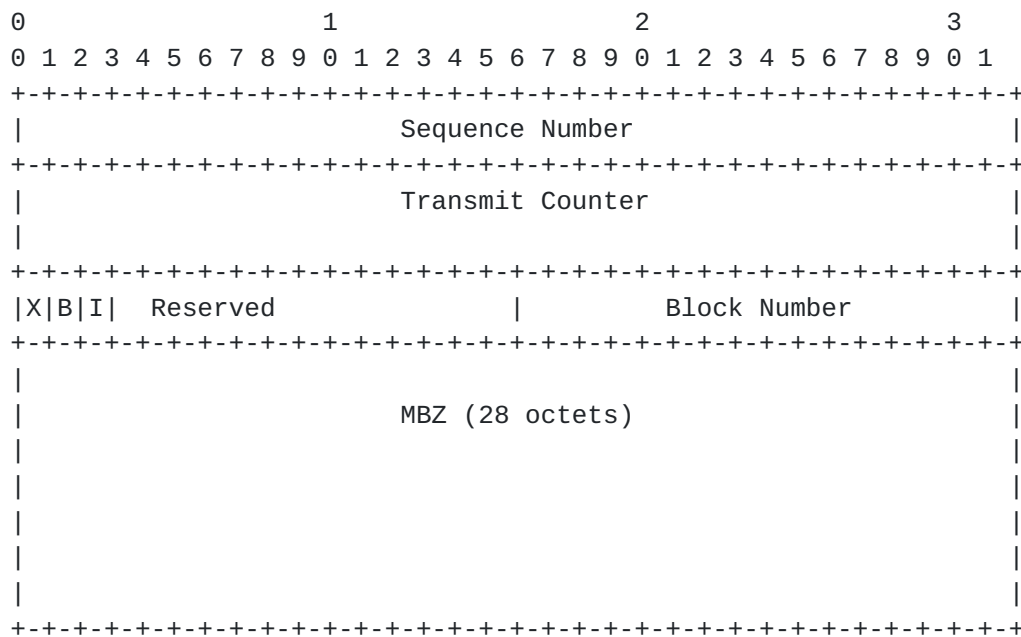


Figure 3C: STAMP LM Probe Query Message Format

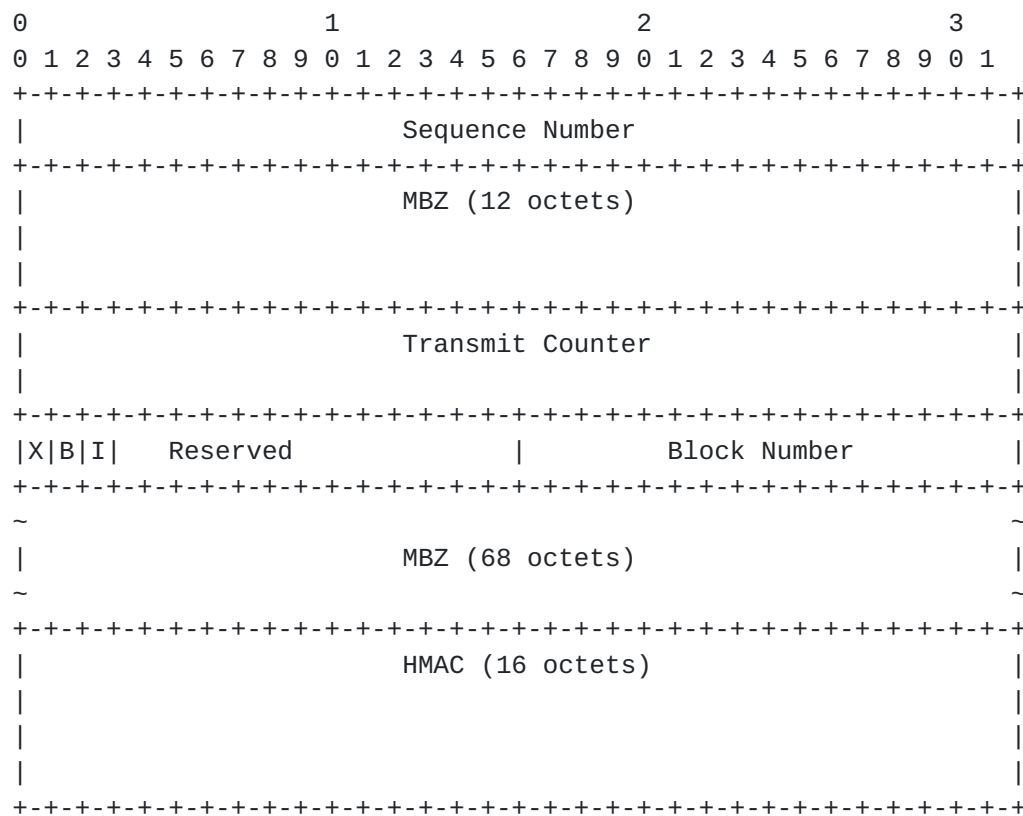


Figure 3D: STAMP LM Probe Query Message Format - Authenticated Mode

Sequence Number (32-bit): As defined in [[RFC5357](#)].

Transmit Counter (64-bit): The number of packets sent by the sender node in the query message and by the responder node in the response message. The counter is always written at the fixed location in the probe query and response messages.

Receive Counter (64-bit): The number of packets received at the responder node. It is written by the responder node in the probe response message.

Sender Counter (64-bit): This is the exact copy of the transmit counter from the received query message. It is written by the responder node in the probe response message.

Sender Sequence Number (32-bit): As defined in [[RFC5357](#)].

Sender TTL: As defined in [[RFC5357](#)].

Flag: The meanings of the Flag bits are:

X: Extended counter format indicator. Indicates the use of extended (64-bit) counter values. Initialized to 1 upon creation (and prior to transmission) of an LM Query and copied from an LM Query to an LM response. Set to 0 when the LM message is transmitted or received over an interface that writes 32-bit counter values.

B: Octet (byte) count. When set to 1, indicates that the Counter 1-4 fields represent octet counts. The octet count applies to all packets within the LM scope, and the octet count of a packet sent or received includes the total length of that packet (but excludes headers, labels, or framing of the channel itself). When set to 0, indicates that the Counter fields represent packet counts.

I: Inferred Mode Loss Measurement: When set to 1, indicates that inferred-mode of loss measurement is used. When set to 0, it indicates direct-mode of loss measurement is used.

Block Number (16-bit): The Loss Measurement using Alternate-Marking method defined in [\[RFC8321\]](#) requires to identify the Block Number (or color) of the traffic counters. The probe query and response messages carry Block Number for the traffic counters for loss measurement. In both probe query and response messages, the counters MUST belong to the same Block Number.

HMAC: The PM probe packet in authenticated mode includes a key Hashed Message Authentication Code (HMAC) ([\[RFC2104\]](#)) hash. Each probe query and response messages are authenticated by adding Sequence Number with Hashed Message Authentication Code (HMAC) TLV. It can use HMAC-SHA-256 truncated to 128 bits (similarly to the use of it in IPsec defined in [\[RFC4868\]](#)); hence the length of the HMAC field is 16 octets.

HMAC uses its own key and the mechanism to distribute the HMAC key is outside the scope of this document.

In authenticated mode, only the sequence number is encrypted, and the other payload fields are sent in clear text. The probe packet MAY include Comp.MBZ (Must Be Zero) variable length field to align the packet on 16 octets boundary.

4.1.2.1. Loss Measurement Authentication Mode

When using the authenticated mode for loss measurement, the matching authentication type (e.g. HMAC-SHA-256) and key are user-configured on both the sender and responder nodes. A separate user-configured destination UDP port is used for the loss measurement in

authentication mode due to the different message format.

4.1.3. Probe Query for SR Links

The probe query message as defined in Figure 1 is sent on the congruent path of the data traffic for Delay measurement. The probe query message as defined in Figure 2 is sent on the congruent path of the data traffic for Loss measurement.

4.1.4. Probe Query for End-to-end Measurement for SR Policy

The performance delay and loss measurement for segment routing is applicable to both SR-MPLS and SRv6 Policies.

4.1.4.1. Probe Query Message for SR-MPLS Policy

The probe query messages for end-to-end performance measurement of an SR-MPLS Policy is sent using its SR-MPLS header containing the MPLS segment list as shown in Figure 4.

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Segment List(1)                | TC | S |          TTL          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
.
.
.
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               Segment List(n)                | TC | S |          TTL          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                               PSID                            | TC | S |          TTL          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Message as shown in Figure 1 for DM or Figure 2 for LM      |
.
+-----+

```

Figure 4: Probe Query Message for SR-MPLS Policy

The Segment List (SL) can be empty to indicate Implicit NULL label case for a single-hop SR Policy.

The Path Segment Identifier (PSID) [[I-D.spring-mpls-path-segment](#)] of the SR-MPLS Policy is used for accounting received traffic on the egress node for loss measurement. The PSID is not required for end-to-end SR Policy delay measurement.

4.1.4.2. Probe Query Message for SRv6 Policy

An SRv6 Policy setup using the SRv6 Segment Routing Header (SRH) and a Segment List as defined in [I-D.6man-segment-routing-header]. The probe query messages for end-to-end performance measurement of an SRv6 Policy is sent using its SRH and Segment List as shown in Figure 5.

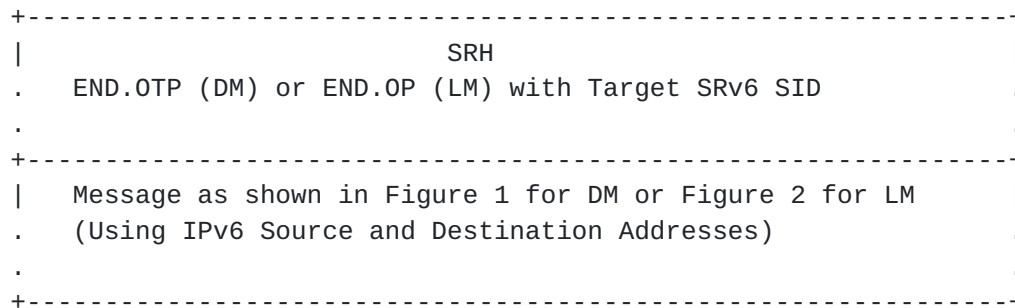
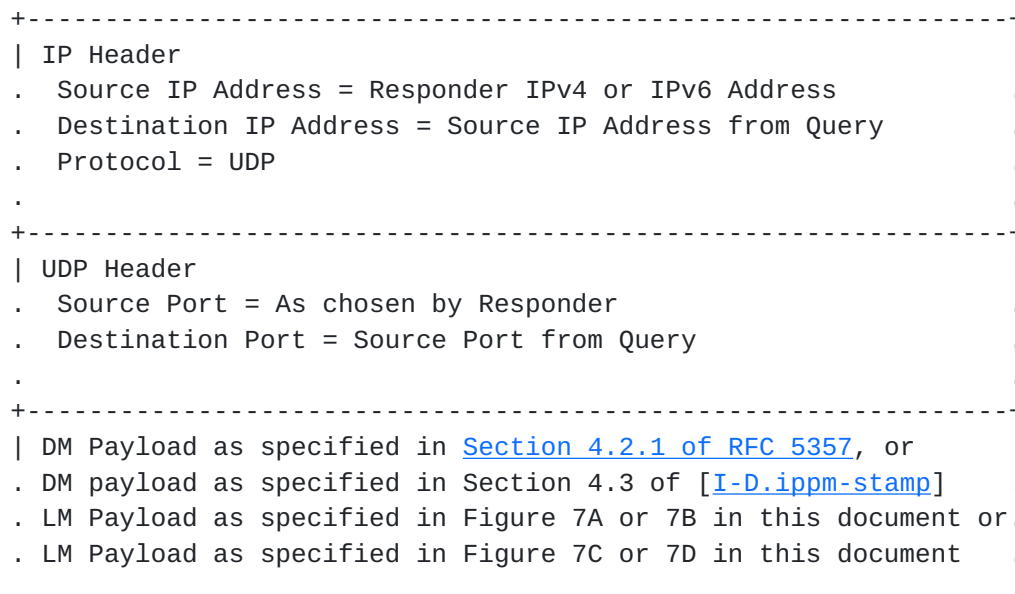


Figure 5: Probe Query Message for SRv6 Policy

For delay measurement of SRv6 Policy using SRH, END function END.OTP [I-D.6man-srv6-oam] is used with the target SRv6 SID to punt probe messages on the target node, as shown in Figure 5. Similarly, for loss measurement of SRv6 Policy, END function END.OP [I-D.6man-srv6-oam] is used with target SRv6 SID to punt probe messages on the target node.

4.2. Probe Response Message

The probe response message is sent using the IP/UDP information from the received probe query message. The content of the probe response message is shown in Figure 6.



+-----+

Figure 6: Probe Response Message

In this document, new probe response message formats are defined for loss measurement as shown in Figure 7A and Figure 7B. The message formats are hardware efficient due to the small size payload and well known locations of the counters. They are also similar to the delay measurement message formats.

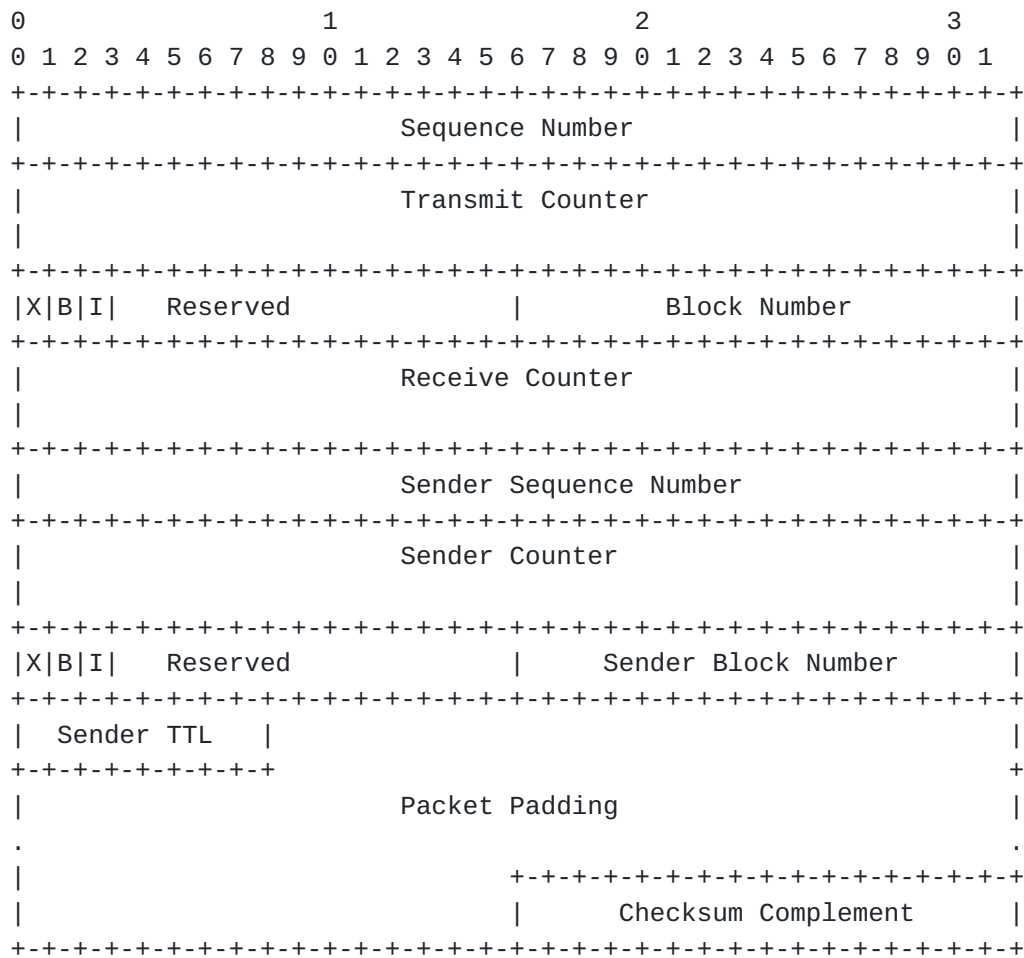



Figure 7A: LM Probe Response Message Format



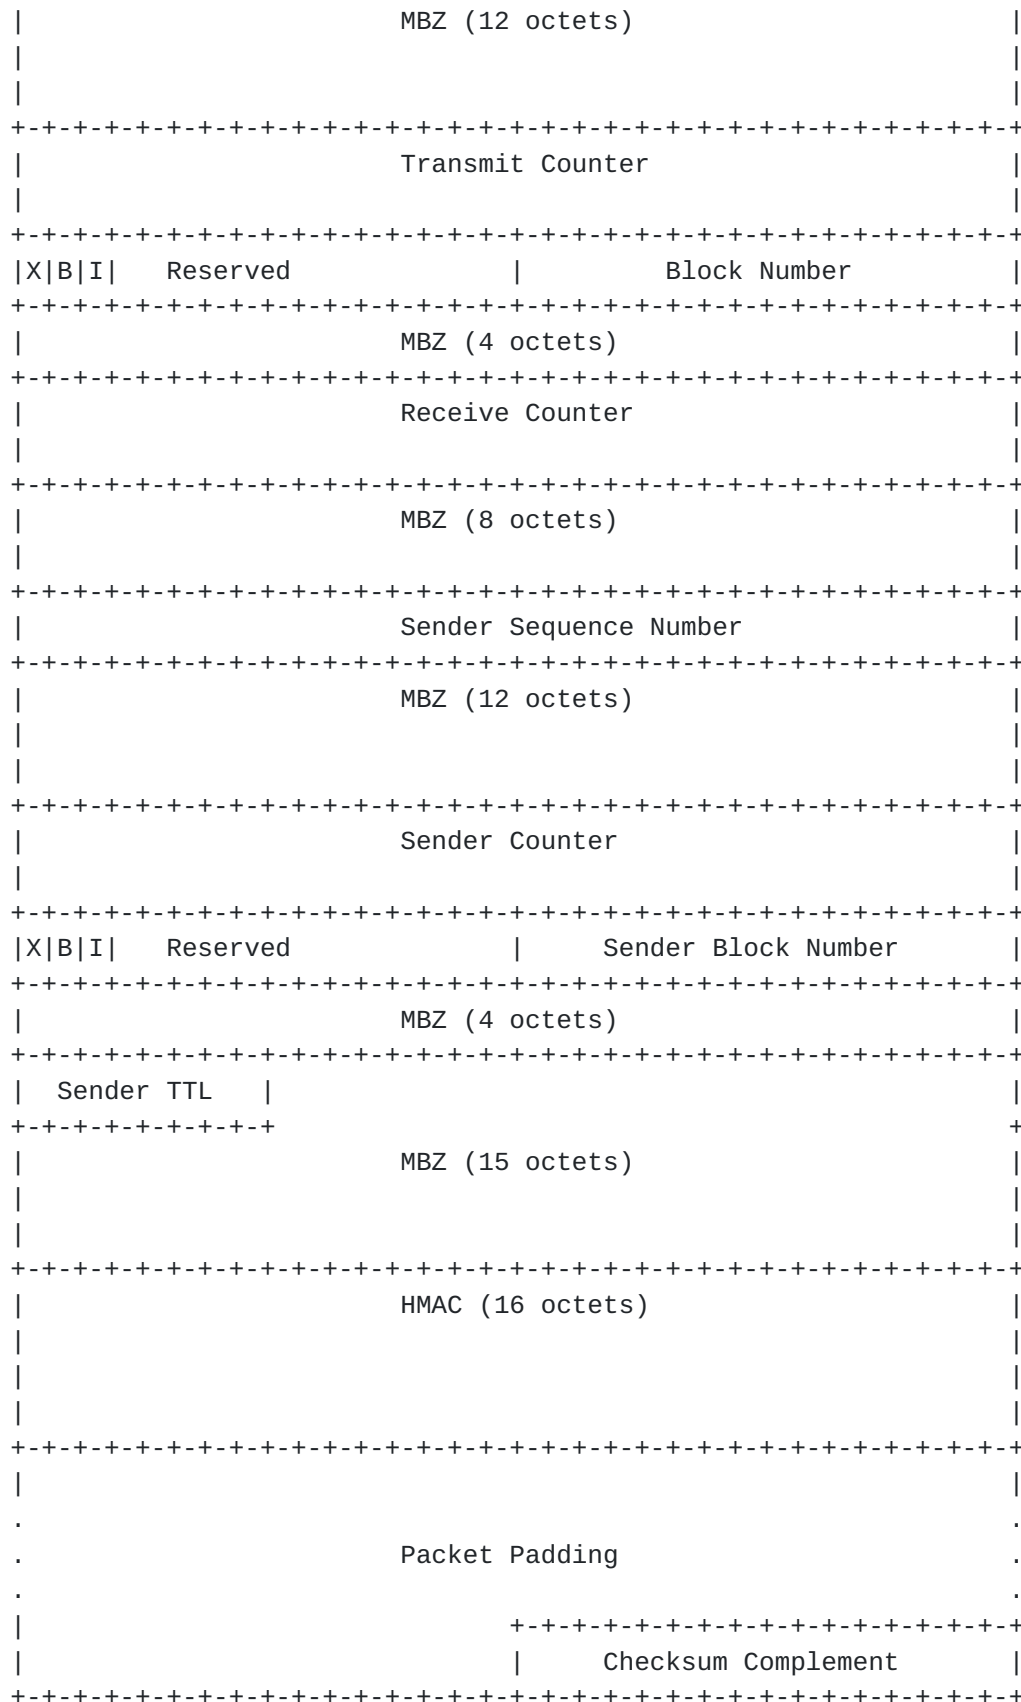


Figure 7B: LM Probe Response Message Format - Authenticated Mode

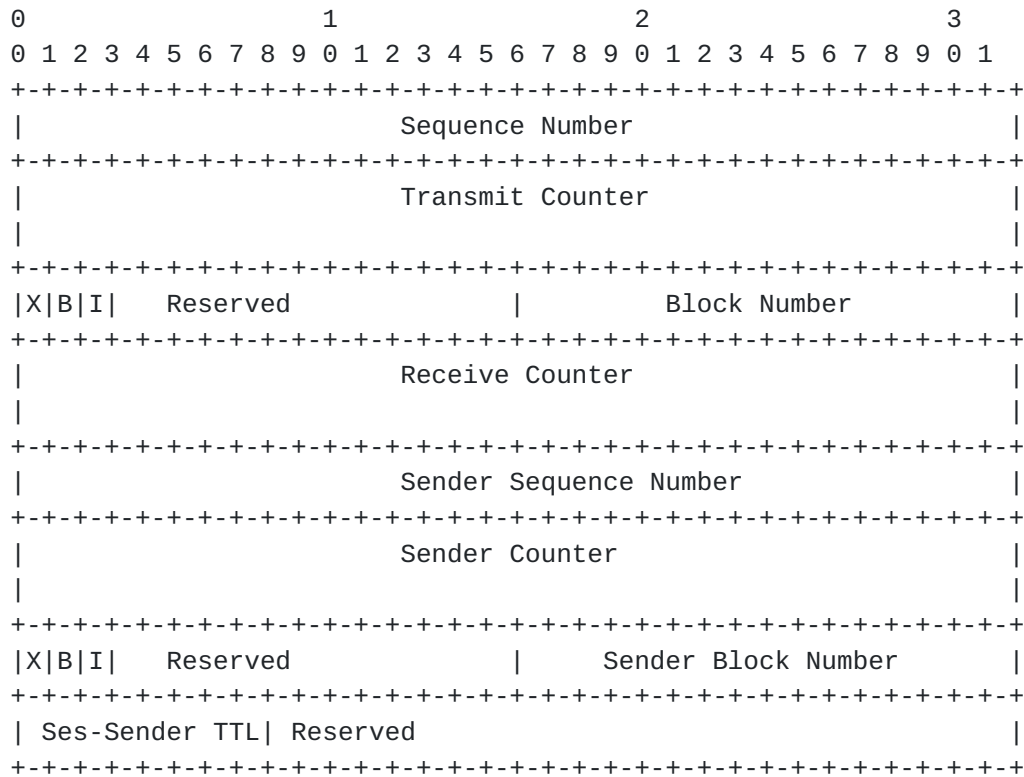
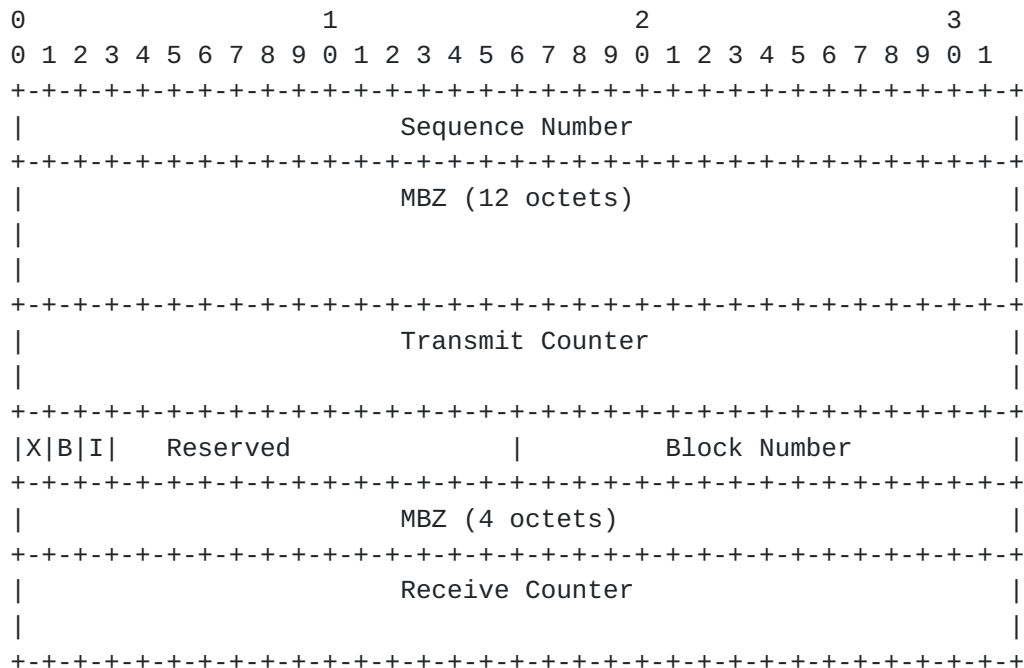


Figure 7C: STAMP LM Probe Response Message Format



For two-way performance measurement, the responder node needs to send the probe response message on a specific reverse path. This way the

destination node does not require any additional state. The sender node can request in the probe query message to the responder node to send a response back on a given reverse path (e.g. co-routed path for two-way measurement).

[I-D.ippm-stamp-option-tlv] defines STAMP probe query messages that can include one or more optional TLVs. New TLV Type (TBA1) is defined in this document for Return Path to carry reverse path for probe response messages (in the payload of the message). The format of the Return Path TLV is shown in Figure 8A and Figure 8B:

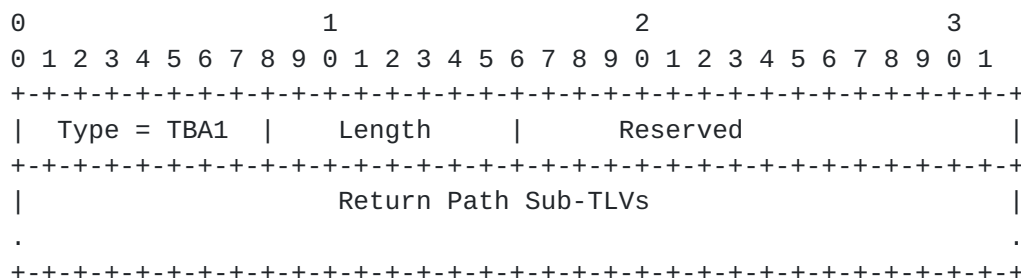


Figure 8A: Return Path TLV

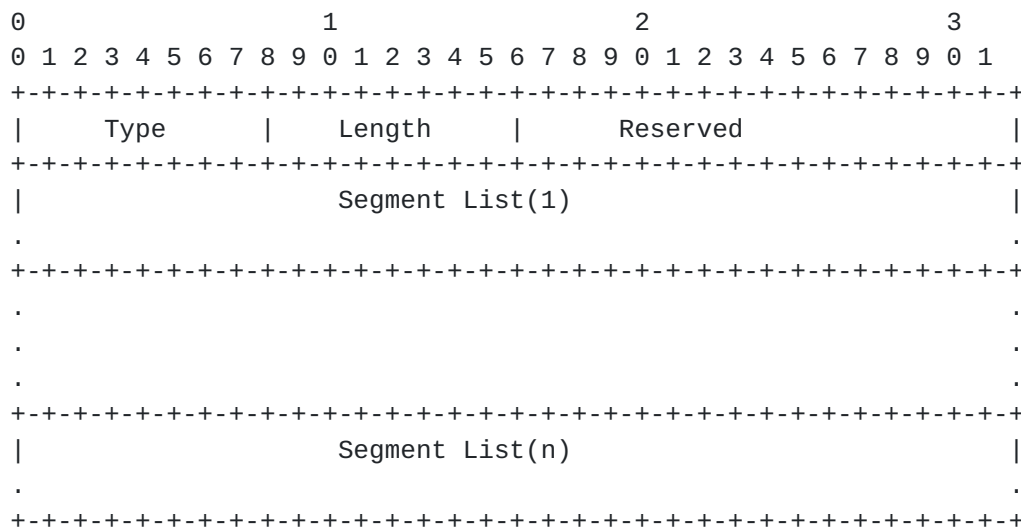


Figure 8B: Segment List Sub-TLV in Return Path TLV

The Segment List Sub-TLV in the Return Path TLV can be one of the following Types:

- o Type (value 1): Respond back on Incoming Interface (Layer-3 and Layer-2) (Segment List is Empty)

- o Type (value 2): SR-MPLS Segment List (Label Stack) of the Reverse SR Path
- o Type (value 3): SR-MPLS Binding SID [[I-D.pce-binding-label-sid](#)] of the Reverse SR Policy
- o Type (value 4): SRv6 Segment List of the Reverse SR Path
- o Type (value 5): SRv6 Binding SID [[I-D.pce-binding-label-sid](#)] of the Reverse SR Policy

The Return Path TLV is optional. The PM sender node MUST only insert one Return Path TLV in the probe query message and the responder node MUST only process the first Return Path TLV in the probe query message and ignore other Return Path TLVs if present. The responder node MUST send probe response message back on the reverse path specified in the Return Path TLV and MUST NOT add Return Path TLV in the probe response message.

[4.2.2.2.](#) Probe Response Message for SR-MPLS Policy

The message content for sending probe response message for two-way end-to-end performance measurement of an SR-MPLS Policy is shown in Figure 9.

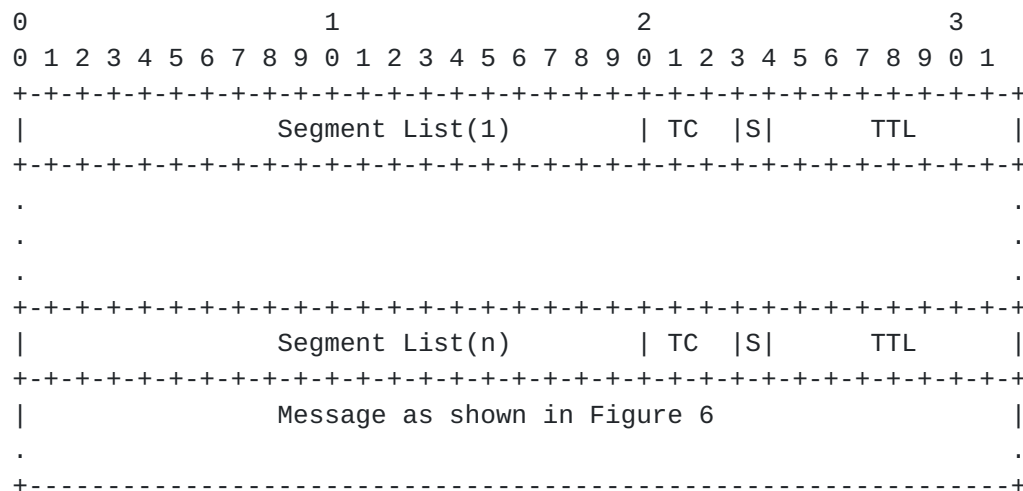


Figure 9: Probe Response Message for SR-MPLS Policy

The Path Segment Identifier (PSID) [[I-D.spring-mpls-path-segment](#)] of the forward SR Policy can be used to find the reverse SR Policy to send the probe response message for two-way measurement of SR Policy.

[4.2.2.3.](#) Probe Response Message for SRv6 Policy

The message content for sending probe response message on the congruent path of the data traffic for two-way end-to-end performance measurement of an SRv6 Policy with SRH is shown in Figure 10.

```

+-----+
|                               SRH                               |
|  END.OTP (DM) or END.OP (LM) with Target SRv6 SID              |
|  .                                                              |
|  .                                                              |
+-----+
|  Message as shown in Figure 6                                  |
|  .  (with IPv6 Source and Destination Addresses)                |
|  .                                                              |
+-----+

```

Figure 10: Probe Response Message for SRv6 Policy

4.2.3. Loopback Measurement Mode

The Loopback measurement mode can be used to measure round-trip delay for a bidirectional SR Path. The IP header of the probe query message contains the destination address equals to the sender address and the source address equals to the responder address. Optionally, the probe query message can carry the reverse path information (e.g. reverse path label stack for SR-MPLS) as part of the SR header. The responder node does not process the PM probe messages and generate response messages.

5. Performance Measurement for P2MP SR Policies

The procedures for delay and loss measurement described in this document for Point-to-Point (P2P) SR Policies [[I-D.spring-segment-routing-policy](#)] are also equally applicable to the Point-to-Multipoint (P2MP) SR Policies [[I-D.spring-sr-p2mp-policy](#)] as following:

- o The sender root node sends probe query messages using the Replication Segment defined in [[I-D.spring-sr-p2mp-policy](#)] for the P2MP SR Policy as shown in Figure 11.
- o Each responder leaf node sends its IP address in the Source Address of the probe response messages. This allows the sender root node to identify the responder leaf nodes of the P2MP SR Policy.
- o The P2MP root node measures the end-to-end delay and loss performance for each P2MP leaf node of the P2MP SR Policy.

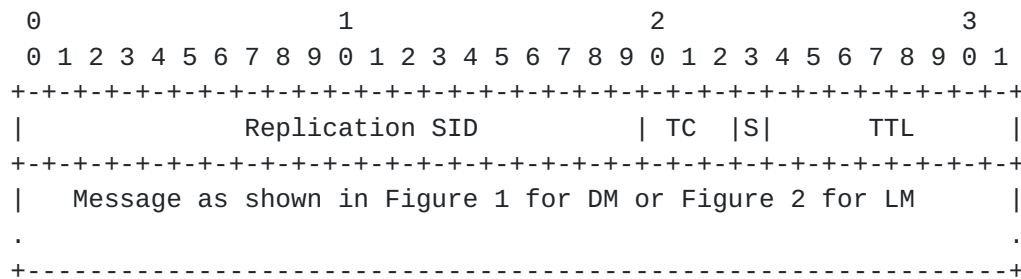


Figure 11: With Replication Segment for SR-MPLS Policy

6. ECMP Support for SR Policies

An SR Policy can have ECMPs between the source and transit nodes, between transit nodes and between transit and destination nodes. Usage of Anycast SID [RFC8402] by an SR Policy can result in ECMP paths via transit nodes part of that Anycast group. The PM probe messages need to be sent to traverse different ECMP paths to measure performance delay of an SR Policy.

Forwarding plane has various hashing functions available to forward packets on specific ECMP paths. The mechanisms described in [RFC8029] and [RFC5884] for handling ECMPs are also applicable to the performance measurement. In the IP header of the PM probe messages, Destination Addresses in 127/8 range for IPv4 or 0:0:0:0:0:FFFF:7F00/104 range for IPv6 can be used to exercise a particular ECMP path. As specified in [RFC6437], 3-tuple of Flow Label, Source Address and Destination Address fields in the IPv6 header can also be used.

7. Additional Message Processing Rules

7.1. TTL Value

The TTL or the Hop Limit field in the IP, MPLS and SRH headers of the probe query messages are set to 255 [RFC5357].

When using the Destination IPv4 Address from the 127/8 range, the TTL in the IPv4 header is set to 1 [RFC8029]. Similarly, when using the Destination IPv6 Address from the 0:0:0:0:0:FFFF:7F00/104 range, the Hop Limit field in the inner IPv6 header is set to 1 whereas in the outer IPv6 header is set to 255.

7.2. Router Alert Option

The Router Alert IP option is not set when using the routable

Destination IP Address in the probe messages.

When using the Destination IPv4 Address from the 127/8 range, to be able to punt probe packets on the responder node, the Router Alert IP Option of value 0x0 [[RFC2113](#)] for IPv4 MAY be added [[RFC8029](#)]. Similarly, when using the Destination IPv6 Address from the 0:0:0:0:0:FFFF:7F00/104 range, the Router Alert IP Option of value 69 [[RFC7506](#)] for IPv6 MAY be added in the destination option. For SRv6 Policy, it is added in the inner IPv6 header.

7.3. UDP Checksum

The Checksum Complement for delay and loss measurement messages follows the procedure defined in [[RFC7820](#)] and can be optionally used with the procedures defined in this document.

For IPv4 and IPv6 probe messages, where the hardware is not capable of re-computing the UDP checksum or adding checksum complement [[RFC7820](#)], the sender node sets the UDP checksum to 0 [[RFC6936](#)] [[RFC8085](#)]. The receiving node bypasses the checksum validation and accepts the packets with UDP checksum of 0 for the UDP port being used for PM.

8. Security Considerations

The performance measurement is intended for deployment in well-managed private and service provider networks. As such, it assumes that a node involved in a measurement operation has previously verified the integrity of the path and the identity of the far-end responder node.

If desired, attacks can be mitigated by performing basic validation and sanity checks, at the sender, of the counter or timestamp fields in received measurement response messages. The minimal state associated with these protocols also limits the extent of measurement disruption that can be caused by a corrupt or invalid message to a single query/response cycle.

Use of HMAC-SHA-256 in the authenticated mode protects the data integrity of the probe messages. SRv6 has HMAC protection authentication defined for SRH [[I-D.6man-segment-routing-header](#)]. Hence, PM probe messages for SRv6 may not need authentication mode. Cryptographic measures may be enhanced by the correct configuration of access-control lists and firewalls.

9. IANA Considerations

IANA is requested to allocate value for the following Return Path TLV Type for [[I-D.ippm-stamp-option-tlv](#)] to be carried in PM probe query messages:

- o Type TBA1: Return Path TLV

IANA is also requested to allocate the values for the following Sub-TLV Types for the Return Path TLV.

- o Type (value 1): Respond back on Incoming Interface (Layer-3 and Layer-2) (Segment List is Empty)
- o Type (value 2): SR-MPLS Segment List (Label Stack) of the Reverse SR Path
- o Type (value 3): SR-MPLS Binding SID [[I-D.pce-binding-label-sid](#)] of the Reverse SR Policy
- o Type (value 4): SRv6 Segment List of the Reverse SR Path
- o Type (value 5): SRv6 Binding SID [[I-D.pce-binding-label-sid](#)] of the Reverse SR Policy

[10.](#) References

[10.1.](#) Normative References

- [RFC768] Postel, J., "User Datagram Protocol", STD 6, [RFC 768](#), August 1980.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.
- [RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", [RFC 4656](#), September 2006.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", [RFC 5357](#), October 2008.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [RFC 8174](#), May 2017.
- [I-D.6man-srv6-oam] Ali, Z., et al., "Operations, Administration, and Maintenance (OAM) in Segment Routing Networks with IPv6 Data plane (SRv6)", [draft-ietf-6man-spring-srv6-oam](#),

work in progress.

[I-D.ippm-stamp] Mirsky, G. et al. "Simple Two-way Active Measurement Protocol", [draft-ietf-ippm-stamp](#), work in progress.

[I-D.ippm-stamp-option-tlv] Mirsky, G., et al., "Simple Two-way Active Measurement Protocol Optional Extensions", [draft-ietf-ippm-stamp-option-tlv](#), work in progress.

10.2. Informative References

- [IEEE1588] IEEE, "1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", March 2008.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), February 1997.
- [RFC2113] Katz, D., "IP Router Alert Option", [RFC 2113](#), February 1997.
- [RFC4868] Kelly, S. and S. Frankel, "Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec", [RFC 4868](#), May 2007.
- [RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", [RFC 5884](#), DOI 10.17487/RFC5884, June 2010.
- [RFC6038] Morton, A. and L. Ciavattone, "Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features", [RFC 6038](#), October, 2010
- [RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", [BCP 165](#), [RFC 6335](#), August 2011.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", [RFC 6437](#), November 2011.
- [RFC6936] Fairhurst, G. and M. Westerlund, "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums", [RFC 6936](#), April 2013.

- [RFC7506] Raza, K., Akiya, N., and C. Pignataro, "IPv6 Router Alert Option for MPLS Operations, Administration, and Maintenance (OAM)", [RFC 7506](#), DOI 10.17487/RFC7506, April 2015, <<http://www.rfc-editor.org/info/rfc7506>>.
- [RFC7820] Mizrahi, T., "UDP Checksum Complement in the One-Way Active Measurement Protocol (OWAMP) and Two-Way Active Measurement Protocol (TWAMP)", [RFC 7820](#), March 2016.
- [RFC8029] Kompella, K., Swallow, G., Pignataro, C., Kumar, N., Aldrin, S. and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", [RFC 8029](#), March 2017.
- [RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", [BCP 145](#), [RFC 8085](#), DOI 10.17487/RFC8085, March 2017, <<http://www.rfc-editor.org/info/rfc8085>>.
- [RFC8186] Mirsky, G., and I. Meilik, "Support of the IEEE 1588 Timestamp Format in a Two-Way Active Measurement Protocol (TWAMP)", [RFC 8186](#), June 2017.
- [RFC8321] Fioccola, G. Ed., "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", [RFC 8321](#), January 2018.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [I-D.spring-segment-routing-policy] Filsfils, C., et al., "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy](#), work in progress.
- [I-D.spring-sr-p2mp-policy] Voyer, D. Ed., et al., "SR Replication Segment for Multi-point Service Delivery", [draft-voyer-spring-sr-replication-segment](#), work in progress.
- [I-D.spring-mpls-path-segment] Cheng, W., et al., "Path Segment in MPLS Based Segment Routing Network", [draft-ietf-spring-mpls-path-segment](#), work in progress.
- [I-D.6man-segment-routing-header] Filsfils, C., et al., "IPv6 Segment Routing Header (SRH)", [draft-ietf-6man-segment-routing-header](#), work in progress.

[I-D.pce-binding-label-sid] Filsfils, C., et al., "Carrying Binding Label/ Segment-ID in PCE-based Networks",
[draft-ietf-pce-binding-label-sid](#), work in progress.

[BBF.TR-390] "Performance Measurement from IP Edge to Customer Equipment using TWAMP Light", BBF TR-390, May 2017.

[I-D.spring-ioam-sr-mpls] Gandhi, R. Ed., et al., "Segment Routing with MPLS Data Plane Encapsulation for In-situ OAM Data",
[draft-gandhi-spring-ioam-sr-mpls](#), work in progress.

[I-D.spring-ioam-srv6] Ali, Z., et al., "Segment Routing Header encapsulation for In-situ OAM Data",
[draft-ali-spring-ioam-srv6](#), work in progress.

Acknowledgments

The authors would like to thank Thierry Couture for the discussions on the use-cases for TWAMP Light in Segment Routing. The authors would also like to thank Greg Mirsky for reviewing this document and providing useful comments and suggestions. Patrick Khordoc and Radu Valceanu, both from Cisco Systems have helped significantly improve the mechanisms defined in this document. The authors would like to acknowledge the earlier work on the loss measurement using TWAMP described in [draft-xiao-ippm-twamp-ext-direct-loss](#).

Authors' Addresses

Rakesh Gandhi (editor)
Cisco Systems, Inc.
Canada
Email: rgandhi@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Email: cfilsfil@cisco.com

Daniel Voyer
Bell Canada
Email: daniel.voyer@bell.ca

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

Bart Janssens

Colt

Email: Bart.Janssens@colt.net